



Learning from African Classroom Pedagogy to Increase Student Engagement in Education Technologies

Judith Uchidiuno; Evelyn Yarzebinski; Emily Keebler; Kenneth Koedinger; Amy Ogan

Human-Computer Interaction Institute, Carnegie Mellon University

Pittsburgh, Pennsylvania, USA

(jio; eey2; ekeebler; kk1u; aeo)@andrew.cmu.edu

ABSTRACT

Tablet-based educational technologies provide a supplement to traditional classroom-based early literacy education, especially in regions with limited schooling resources. Prior work has probed how children generally interact with and learn from these technologies, however, there is limited research on *student engagement* with applications that utilize valuable input techniques such as automatic handwriting and speech recognition. In our study, we designed and field-tested early literacy speech and handwriting recognition applications with the primary aim of maximizing student engagement. We designed the applications based on prior research insights and classroom observations from our target population and field-tested the applications with 283 children living in rural Tanzania. We found that observing a small set of classrooms can produce design insights that increase engagement on tablet-based learning systems on a much larger scale. We also demonstrate the importance of domain familiarity in students' choice to persist through activities while learning with technology.

CCS CONCEPTS

- Applied computing~Computer-assisted instruction

KEYWORDS

Early literacy; educational technologies; Sub-Saharan Africa; speech recognition; handwriting recognition

1 Introduction

There have been several initiatives dedicated to designing and deploying educational technologies to improve early literacy skills in Sub-Saharan African regions [4,26,39]. These efforts have been made in an attempt to reduce the educational disparities between rural and urban areas arising from little to no access to traditional schooling infrastructure

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from Permissions@acm.org.

COMPASS '19, July 3–5, 2019, Accra, Ghana
© 2019 Association for Computing Machinery.
ACM ISBN 978-1-4503-6714-1/19/07...\$15.00
<https://doi.org/10.1145/3314344.3332501>

and resources such as skilled teachers, funding for physical structures, and other basic necessities like textbooks [30,31]. Recent estimates by the United Nations International Children's Emergency Fund (UNICEF) indicate that of the over 57 million children who are currently out of school, over half of them live in sub-Saharan Africa [33]. Other estimates show that sub-Saharan Africa is home to the highest number of school children who do not pass minimum proficiency thresholds in both mathematics and reading [37].

The affordability, scalability, and ubiquity of mobile devices in developing regions have made them a popular solution in the attempt to close these proficiency gaps. Technologies that have shown potential to produce learning gains include traditional desktop computers [26], tablets [4,39], and even feature phones using interactive voice response technologies [24]. Most of these efforts provide valuable insights about children's engagement with such devices using direct manipulation techniques such as tapping; however, they do not take advantage of more advanced features of such technologies such as speech and handwriting recognition – features that can provide on time, personalized feedback and scaffolding for children's reading and writing instruction. In addition, these studies provide no insights on factors that motivate children to engage with and learn from these technologies, especially when deployed without scaffolding from teachers or knowledgeable adults.

It is especially important to provide educational solutions that encourage children to engage with such technologies independently, as the presence of a knowledgeable adult is not guaranteed. Studies have shown that families in rural contexts often depend on children's labor to survive and may value other forms of work over formal education [19]. Therefore, designing applications that are highly engaging is especially important – children have to choose to engage with these technologies despite their other time commitments and limited family support. Given a lack of published research studies that explore factors that affect children's engagement with learning technologies specifically in Africa, we investigated student engagement studies conducted in other regions in the world, as well as conducted classroom observations in the local African region to inform the design of reading (with automatic speech recognition), and writing (with automatic handwriting recognition) tablet-based applications targeted at Swahili-speaking K-2 students living in rural regions in Tanzania.

The more time children spend on these applications, the more learning can potentially occur, therefore, we use completion rates (student persistence) on individual activities as a measure of student engagement.

In this paper, we describe our efforts to *understand* highly engaging pedagogical techniques by observing local classrooms. We *incorporate* those techniques into the design of our reading and writing tablet applications, and *evaluate* their effectiveness with out-of-school students learning primarily with tablets. We observed classroom teaching in one region (Region-1) over a 2-week period, incorporated the insights learned into the design of our reading and writing applications, deployed our applications to 283 students in a different but culturally and economically similar region (Region-2) over a 6-month period, and report on the effects of incorporating the learned insights on student engagement with the tablet applications. We explore student engagement by reviewing completion rates in different learning activities, and discuss the features that influence student persistence. We make the following valuable contributions to the learning science and human-computer interaction research communities:

- We provide evidence that incorporating pedagogical techniques observed from a small set of local classrooms into the design of education technologies can increase student engagement in a much larger, technology-based, non-traditional learning environment without teacher scaffolding.
- Contrary to some prior research studies, we provide evidence that activity difficulty does not always predict student engagement – other factors such as domain familiarity may lead to high levels of engagement even with difficult activities.

2 Background

2.1 Early Literacy Education Technologies in Developing Contexts

The rising number of out-of-school children all over the world has led several organizations to attempt to tackle the issue using interventions that scale much faster than traditional schooling. Initiatives that use educational technologies as a means of closing educational disparities have become increasingly popular e.g., [3,4,18,27]. There have been few such initiatives conducted in African regions specifically – some directed at augmenting inadequate traditional schooling, while others focus on supporting out-of-school students. The One Laptop Per Child (OLPC) program provided primary school students from several African countries e.g., Tanzania, Ghana etc. with individual low-cost laptops to support their learning in school. Unfortunately, most of its success was hindered by the lack of infrastructure, support staff, and adequate security for the devices. For out-of-school children, Breazeal et al. deployed

a tablet-based educational application in two villages in Ethiopia, and found that even in out-of-school contexts, over half learned more than 50% of the vocabulary provided [4]. Furthermore, our research team has worked with rural communities in Tanzania over the last two years, reporting on how children interact with tablet-based education systems in both home and school settings, and providing insights on the types of struggles children face in the learning process [39]. In this study, we continue our work in these communities by gathering insights that inform the design of reading and writing applications that maximize student engagement in rural, Swahili-speaking, non-traditional tablet learning contexts in Tanzania.

2.2 The Importance of Teaching Early Reading and Writing Skills

It is widely accepted that teaching reading skills (e.g., letter shapes and sounds, words, sentences, comprehension) is a critical component of any formal education program [35]. Reading allows students to learn “about other people, about history and social studies, the language arts, science, mathematics, and the other content subjects that must be mastered in school” [23]. It allows students to navigate the world as global citizens, to obtain economically lucrative jobs outside of their local communities, and to cope and adjust to new information as the world evolves [29].

Handwriting training is also a very important component of any early childhood literacy intervention. It improves fine motor skills and stimulates areas of the brain that control self-regulation and executive function [7]. Children who are learning to write, especially copying characters, must self-regulate, demonstrate impulse control and attentional flexibility, and utilize working memory [8]. Teaching handwriting skills (including tracing and copying letters) in early literacy education has been associated with better academic performance in different domain areas including math and reading among K-1 students [15,22,36].

2.3 Student Engagement in Education Technologies

Individuals choose to be engaged when activities are “interesting, socially useful or personally enjoyable to them, and within their zone of proximal development. It involves self-direction, interaction, emotion, choice and a sense of competence” [32]. There are several indicators of children’s engagement with learning materials, such as moving physically closer to the materials when responding to questions [6], persisting through activities by pointing and page turning [9], and expressing both enthusiasm through smiling and laughing and compliance such as through timely responding, staying seated, and following directions [28]. Children use a variety of multisensory behaviors such as looking, touching, listening, and even facial expressions and verbal and non-verbal communication to express their interest in learning materials [28].

In non-African regions, student engagement with learning technologies has been measured using different metrics such as duration of voluntary play [20], number of game levels attempted [2], etc. Researchers have controlled for several factors including the addition of background music, varying game difficulty, celebratory animations, and optional bonus challenges to uncover factors that affect student engagement. Andersen et al. found that adding background music to educational games, had no significant effect on student engagement [2], but they did not explore the use of other types of music e.g., nursery rhymes, or teaching domain content using songs. This distinction is very important because singing and songwriting have been shown to be correlated with better academic performance [25], improved learning of letter names and sounds [11], phonemic awareness [12], and improved word identification [34]. Andersen et al. also found that celebratory animations showed increased engagement, but the presence of optional challenges showed the reverse effect. Other studies have explored the effect of activity difficulty on student engagement, with contradicting results. While some research studies show that learners always choose games with a moderate level of difficulty [1,13], Lomas et al. found that in their learning system, learners showed the highest levels of engagement with the easiest games [21]. These mixed results imply that, depending on context, activity difficulty may not always predict student engagement with learning applications.

We investigate the following research questions:

1. Do pedagogical techniques that seem engaging to children during traditional classroom instruction also engage children in tablet-based instruction?
2. Other than activity difficulty, what factors explain student engagement with tablet-based learning applications?

3 Methodology

This study is part of a larger effort to design Android tablet-based early literacy and numeracy applications designed for children in Sub-Saharan Africa. The results reported in this paper are a subset of our ongoing work with various communities in the region. This study was conducted in partnership with a rural Swahili-speaking village in the Northwestern part of Tanzania (Region-1) and a global organization focused on reducing education disparities in rural villages in Tanzania (these rural villages will be referred to as Region-2). Members of our research team have conducted research in Region-1 over the last three years designing and deploying early literacy applications to supplement traditional classroom learning and are familiar with the local conditions and context.

At the time of our last visit, Region-1 was quite limited in physical and technical infrastructure. Power was somewhat consistent; however, after one night of heavy rainfall, there

was no power until the research team left the region. The closest ATM or medical clinic was at least an hour away, and the first water well was constructed at the time of data collection. Three different mobile network providers serviced the area, and we purchased SIM cards from all three networks to determine the network with the fastest data speeds. The fastest network speeds in the village was EDGE (Enhanced Data rates for GSM Evolution) with download speeds ranging from 120Kbps to 384Kbps. We did not observe any tablets in the region either in home, school, or public settings throughout the duration of our stay. The predominant occupation of parents in Region-1 is farming. The region has two primary schools (grades K-7): an English-instruction private school that costs \$600 a year and a Swahili-instruction public school with free tuition but for which parents are still responsible for certain fees, uniforms, and school supplies.

Prior to the current iteration of our reading and writing applications, our team visited K-2 classrooms in both schools in Region-1 daily for 2 weeks, observing teaching methodologies, student and teacher engagement, and student interactions in the classroom. Our aim was to gain insights on engaging teaching practices and incorporate them into the design of our tablet applications. We obtained permission from the village council, school principal, and classroom teachers, and students were informed of our presence and our research goals prior to conducting the classroom observations. All classroom observation notes were handwritten for further analysis, and later discussed with the teachers and other village residents to provide additional context to our observations. *For this study, all classroom observations were conducted in the two schools in Region-1.*

Finally, we partnered with a global organization as part of a larger effort to provide education to children with no schooling in Region-2, a culturally and geographically similar area to Region-1. As in Region-1, the primary parent occupation in Region-2 is farming (79%), with only 8% of parents having above basic primary (elementary) school education. The global organization distributed individual tablets to school age students in the target region after receiving consent from their parents to participate in the research study. Information about the study was disseminated from the village council, and parents were asked to bring their children for registration if they were interested. Verbal consent was obtained from the parents to allow their children participate in the study and to allow the gathering of de-identified data for further analysis from the tablets. Our learning application was deployed to 283 children ages 5-12 in Region-2. About 39% of the participants were girls, and only 11% of all participants were enrolled in school as reported by the parents. The main reasons parents cited for not enrolling their children in school included the distance from their home, the cost of school, and health concerns. Only 2% of the children reported ever using a smartphone (and much less a tablet). Although 70% of the parents reported owning mobile phones, only 3% of the

parents reported ever using a smartphone (half of that for tablet usage). *For this study, all usage logs, performance and engagement data will be reported based on analysis of data from Region-2.*

3.1 Learning Application

Our application focuses on identifying, writing, and speaking in the following domains: literacy (letter names, syllable sounds, words, sentences); mathematics (numbers, basic arithmetic, word problems, number comparisons, and missing numbers); and reading (curated from <http://www.africanstorybook.org/>). System interactions involve screen tapping, writing, and speaking. All system content was recorded in Swahili by a Kenyan Professor of Swahili, who has taught Swahili from kindergarten to university levels. Additional video tutorials were recorded in Swahili by a Tanzanian instructor of Swahili, who grew up in a region about 10 hours away from Region-1 and Region-2. Although our system contains several applications, in this paper we focus on its reading and handwriting applications. For all applications, students can choose to exit an activity before it is complete using the blue arrow button shown at the top-left of the screen in Figure 1.

3.1.1. Speech Recognition Application Features

Our Automatic Speech Recognizer (ASR) was adapted from a lightweight speech recognition engine, specifically tuned for handheld and mobile devices (PocketSphinx [14]). This tool has been used in prior research to enable speech recognition in mobile games for children in rural India [17]. It was modified for use as a Swahili ASR by including a Swahili pronunciation dictionary, and a pronunciation synthesis parser program, while maintaining the language model generator, the phoneme set, and an acoustic model used in a previously developed English ASR. Figure 1 shows a screenshot of a participant reading a story, with words highlighted to provide correctness feedback.



Figure 1: Screenshot of Pilot Student Reading a Story with Right/Wrong Feedback and Pronunciation Support Button

Students can tap on the current word (underlined), or the green audio button to hear the system pronounce the highlighted word. When the application detects a few seconds of silence, the instructions are repeated, followed by a finger hovering over the audio button to remind them where to tap for help. The current version of our reading application

(including changes incorporated after the classroom observations) consists of the following modes:

- **HEAR:** Students listen to text (including letters, numbers, words, storybooks, and songs), without the need to provide a response.
- **PARROT:** This mode is identical to rote repetition. Students hear and see some text (sometimes with an accompanying picture) and are asked to repeat after the application.
- **ECHO:** This mode is the inverse of “parrot” mode above. Students are shown only text and are asked to read what they see. The application provides the correct pronunciation to the text on the screen after the student response regardless of correctness.
- **REVEAL:** Students are shown a picture and are instructed to say what they see. This mode is currently used to assess students’ ability to say a number by looking at a picture of its numeral. The application provides the correct answer after the student response regardless of correctness.
- **READ:** Students are shown text and are asked to read what they see. Unlike the other modes, the application does not provide the correct pronunciation to each word – words are highlighted in different colors to provide correctness feedback.

3.1.2 Handwriting Recognition Application Features

One of the most popular toolkits employed by handwriting recognition researchers, and mobile application developers is Lipi toolkit (<http://lipitk.sourceforge.net>) [40]. It allows for shape and word recognition, and individual character and string recognition. It has been used for developing handwriting applications on android mobile devices, and has been validated for use in recognizing handwriting of children, including those with dyslexia [16]. For our writing application, we modified this toolkit by adding features to score the similarity of the written character to the expected character and biasing it to prefer the expected character. Students provide input by writing on the screen using one finger.

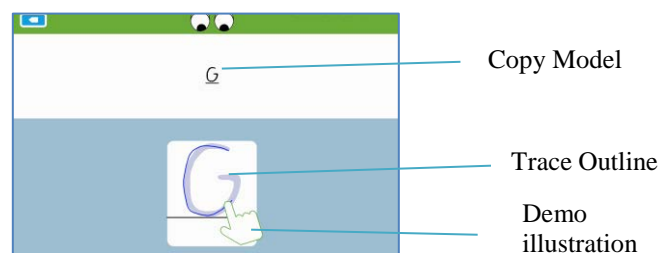


Figure 2: Writing Application – TRACE mode includes copy model and trace outline; COPY mode includes only copy model; DICTATION includes only an audio button to replay target; all modes begin with a demo of the target activity.

Before each activity begins, students are shown a demo with a finger writing the response to an example problem (Figure

2). The current version of our writing application consists of these modes:

- **TRACE:** Students are shown an outline of a letter/syllable/number and are instructed to write over the outline. After three incorrect attempts at a given question, students are advanced to the next question.
- **COPY:** Students see the text to copy. Each subsequent attempt provides a repeat of the instructions or a temporary trace outline. By the fifth incorrect attempt, a trace outline remains in the writing box for scaffolding, and if a wrong response is still provided, students are advanced to the next question.
- **DICTATION:** Students hear a word without any text on the screen and are instructed to write what they hear. The scaffolding provided is similar to ‘COPY’ mode.

3.2 Participants’ Prior Knowledge

Prior to the start of the study, students in both regions were given a similar baseline test to assess their knowledge on various early literacy and numeracy topics. Table 1 and 2 shows the average scores obtained by Region-1 and Region-2 students in the different grades, across the different topics. As expected, the children in Region-1 scored better on the baseline test as they have more years of formal schooling.

Topic	No School N=21 (age 3-6)	Grade K N = 104 (age 4-10)	Grade 1 N = 72 (age 5-9)	Grade 2 N = 59 (age 6-13)
Letter/Syl ID	1%	18%	63%	88%
Word ID	0%	13%	52%	86%
Oral Reading	0%	4%	23%	34%
Listen Comp.	38%	67%	87%	87%
Sentence Cop	0%	11%	58%	86%
Number ID	0%	19%	45%	83%
Basic Add. and Sub.	0%	10%	33%	75%
Word Problems	9%	14%	39%	82%
Number Writing	0%	49%	88%	99%

Table 1: Region-1 Average Pretest Scores

Topic	No School N=214 (age 7-12)	Grade K N = 28 (age 7-11)	Grade 1 N = 21 (age 9-11)	Grade 2 N = 15 (age 9-11)	Grade 3 N = 5 (age 9-11)
Letter/Syl ID	0%	1%	5%	6%	2%
Word ID	0%	0%	4%	6%	6%
Oral Reading	0%	0%	6%	7%	3%
Listen Comp.	39%	44%	65%	57%	56%
Number ID	8%	18%	24%	23%	34%
Basic Add. and Sub.	3%	10%	14%	8%	9%
Word Problems	9%	11%	23%	17%	29%

Table 2: Region-2 Average Pretest Scores – there were no sentence completion, or number writing tasks on the test.

3.3 Data Collection and Analysis

3.3.1 Classroom Observations (Region-1):

We analyzed data from 16 classroom observation sessions (8 from the public school; 8 from the private school), including handwritten observations and photographs of the school and classroom artifacts. Next, we identified emergent themes

related to recognition, reading, and writing instruction following a grounded theory approach [5]. Four themes of instructional pedagogy emerged: 1) the central role of music for all kinds of instruction, 2) situating learning using real-world pictures and familiar contexts, 3) regular use of rote repetition to support knowledge retrieval, and 4) emphasis on syllables (rather than letters) as the unit of reading. Next, two team members reviewed and categorized all notes into these themes, discussed any areas of disagreement, and re-categorized as necessary. Finally, we triangulated all observations related to these themes with other information captured by researchers in the field: debrief recordings, follow-up interviews with teachers, and photographs captured on site to ensure that all evidence were mutually supportive. We worked with native Tanzanian Swahili speakers to help translate interactions, as well as provide insights on any cultural underpinnings.

3.3.2 Application Performance Logs (Region-2):

We analyzed student performance logs, focusing on interactions with the reading and writing applications. Data was collected over a 6-month period (June – November 2018). Student performance was computed as the ratio of student correct attempts to the total number completed attempts in an activity. We operationalized student engagement as student completion rates per activity type (ratio of completed questions to total number of questions in an activity). We ran basic descriptive statistics to reveal high-level differences across the activity types, and for further analysis performed linear regressions and paired t-tests to measure the effects of our annotated features on student engagement on the activities.

We aggregated activities into types by their tutor (read vs write), tutor mode (reading application: hear, parrot, echo, reveal, read, and writing application: trace, copy, and dictation), and domain content (numbers, letters, syllables, regular and nonsense words, missing letters, sentence, etc.). In this research study, we use activity completion rates as a metric to measure engagement, defined as the average proportion of questions a student completes per activity type e.g., if a student completes 2/10 questions in an activity, their completion rate is 20%.

Our curriculum promotes students to harder activities as they achieve mastery. We analyzed data from a total of 202,236 activities attempted by the students – only 42% of those activities were fully completed by the students. Given these numbers, it is crucial to investigate activity features that led to increased completion rates, allowing students to progress through the curriculum at an acceptable rate, especially since motivation and support from teachers or knowledgeable adults is not guaranteed outside of school settings.

4 Findings

4.1 Student progress with our reading and writing applications

Although our study focused on understanding and increasing student engagement, we begin our analysis by presenting student progress using our system over the 6-month period. Figures 3 and 4 show the comparisons between student achievement levels at the beginning and end of the study.

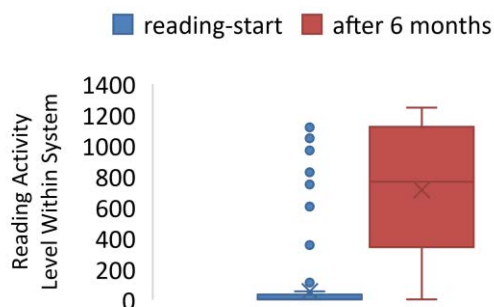


Figure 3: Students' reading levels before (median = level 5 – rote repetition vowels) vs after testing period (median = level 766 – reading of 3-letter syllables); Average use of 6 sessions per month.

Students' median beginning reading level was rote vowel repetition with exposure to vowels and their sounds. Figure 3 shows that 8 outlier students began at above-average levels. Based on the pretest scores, most students in Region-2 could not identify a single letter at the beginning of the study. At the end of the 6-month period, with an average of 6 usage sessions per month, most students were independently reading 3 letter syllables, mastering all 24 letters, and most 2 letter syllables. This increase in student reading levels was statistically significant (*paired t – test: t = -26.689, df = 242, p < 0.01*).

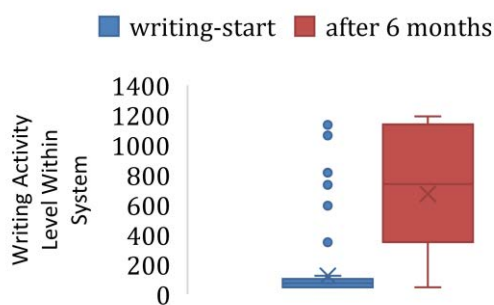


Figure 4: Students writing levels before (median = level 77 – tracing of numbers from 1-4) vs after testing period (median = level 735 – copying of 2-letter syllables); Average use of 6 sessions per month.

For writing, the median beginning level was tracing the vowels. There were 6 outlier students who began at writing levels above the typical student. By the end of the 6-month testing period, with an average usage of 6 usage sessions per month, the median writing level was copying of 2-letter syllables. Students had mastered tracing, copying, and independently writing all numbers from 1-10 and all 24

letters. This difference in student writing levels was statistically significant (*paired t – test: t = -20.533, df = 232, p < 0.01*).

To understand the factors that contributed to these learning gains, the rest of our findings are organized by providing descriptions of our classroom observation for each theme, their design implications, followed by a quantitative analysis of the effects of those design choices on student engagement. We also provide an analysis on the effects of activity difficulty on student engagement.

4.2 The central role of music for all kinds of instruction

4.2.1. Classroom Observation (Region-1)

In the public school, the most prevalent way of teaching students in all the grades observed was through music. In Kindergarten, the teacher used music to teach the children to introduce themselves, to learn about their country's history, to learn about wildlife, and even to acknowledge their peers who answered questions correctly. In grades 1 and 2, music was used at the beginning of each lesson to start the lecture (e.g., a song about the importance of reading), followed by several songs covering the topic of the day (e.g., a song about vowels or numbers), before commencing with the lesson. Most of the songs included an associated dance routine that the children had mastered, and the children seemed highly engaged and excited while singing these songs.

4.2.2. Design Implications

As a result of these observations, we added two types of songs to our reading tablet application to increase student engagement: songs covering specific domain content (e.g., numbers and letters) illustrated with the text of the content, and other songs, similar to nursery rhymes, illustrated with pictures and the text of the songs. All songs were recorded, and all pictures were vetted by a native Swahili speaker born and raised in Tanzania to ensure accurate translation and that the objects matched students' experience (e.g., toilets look very different between urban and rural Tanzanian regions).

4.2.3. Effect on Student Engagement in Tablet-Based Education (Region-2)

We hypothesized that teaching similar content using songs will be more engaging than narrating the same content (or interacting with students using rote repetition) in our learning application as well. We compared students' engagement with learning numbers from 1-10, learning lower case letters, and learning about upper case letters across two teaching methods: songs versus rote repetition activities. Results from a paired t-test shows that teaching with songs was more engaging across all the content types explored:

- Learning numbers from 1-10 (*mean completion = 75% vs 66%; SD = 20% vs 22%; t = 5.0476, df = 150, p < 0.01*)

- Learning lower case letters (*mean completion* = 79% *vs* 55%; *SD* = 20% *vs* 21%; *t* = 13.699, *df* = 207, *p* < 0.01), and
- Learning upper case letters (*mean completion* = 80% *vs* 51%; *SD* = 18% *vs* 18%; *t* = 18.856, *df* = 220, *p* < 0.01).

Andersen et. al [2] compare similar activities with or without music in the background. While we do not have a direct comparison to this variation, we show that teaching children the exact same content using a song is more engaging than teaching them by rote repetition only.

We included two types of songs in our application: local nursery rhymes annotated with photographs and text, and content songs (e.g., the ABC song) annotated with the text of the song. Although all songs were engaging, nursery rhymes annotated with real-world photographs were significantly more engaging than content-area songs (e.g., songs about letters and numbers) annotated with just text (*mean completion* = 81% *vs* 76%; *SD* = 14% *vs* 11%; *t* = -14.042, *df* = 20713, *p* < 0.01). It is worth an exploration to incorporate more real-world pictures in the content specific songs containing domain content, to determine if it leads to students persisting even longer through the content song activities.

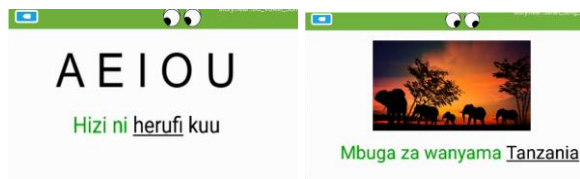


Figure 5: L: Content Songs; R: Nursery Rhymes

4.3 Learning situated using real world pictures and contexts

4.3.1 Classroom Observation (Region-1)

All the learning posters on the classroom walls, and non-textbook materials found in both schools were either real-life photos (rather than cartoon characters) or hand-drawn pictures of relatable concepts. These posters were often used as teaching aids for topics such as types of animals, parts of the body, etc. As part of the class instruction, teachers typically pointed at the items on the posters and asked the children to identify it, before providing correctness feedback.

4.3.2. Design Implications

As a result of our observations, we incorporated several word reading activities with familiar words, including animals, common village objects, family, food, etc. Each activity was accompanied by a relevant photograph (either sourced from an open source image library or created specifically for the application), as well as the text for the target word. Children had opportunities to practice the same words without the accompanying pictures in harder activities.



Figure 6: Familiar words annotated with photographs

4.3.3. Effect on Student Engagement in Tablet-Based Education (Region-2)

We investigated whether students were more engaged with reading familiar words or high-frequency words at their reading level. In our learning application, high frequency words were curated from the African Storybook Project Database, and difficulty was based on word length. Common words were curated based on school observations, the village context (Region-1), and input from adults familiar with the cultural context.

Students showed significantly *less* engagement with reading common words than reading level-appropriate high frequency words (*mean completion* = 71% *vs* 77%; *SD* = 16% *vs* 14%; *t* = -6.6619, *df* = 759.66, *p* < 0.01). Given this unexpected finding, we investigated whether this result was the same across all application types. Although common words were less engaging for both reading activities (PARROT and ECHO), common words were *more* engaging than familiar words in the word-writing task. We hypothesize the presenting children with familiar words narration-style might have been too easy for them (causing possible boredom), however, using those words for a cognitively harder task such as writing provided enough familiarity to persist through the task over other word-writing activities.

4.4 Regular rote repetition to support knowledge retrieval

4.4.1 Classroom Observation (Region-1)

The use of rote repetition was quite common in both schools. The teachers often went through several minutes of writing things on the chalkboard, saying each item while pointing and asking the children to repeat after them. Usually they followed this initial repetition activity with asking the children to identify the items they pointed to without providing help to assess their mastery of the content. In the private school, students often followed along, and when they struggled, relied on one student to provide an answer first before following in chorus. This behavior sometimes went unnoticed by the teacher, except when the whole class provided the wrong answer. The students seemed fairly engaged in this activity, although some more than others. In the public school, the large class sizes made rote repetition a very common in-class activity as well. Even when students were called up to the front, they were asked to answer a question for the whole class to repeat, rather than calling on students for individual responses. Only the students sitting in

front of the public school classes attempted to provide correct responses.

4.4.2 Design Implications

The PARROT mode of our reading application is most similar to rote repetition, and the ECHO/REVEAL/READ modes are most similar to in class activities where the teacher highlights an item on the chalkboard and elicits student response. While these methods had varying levels of engagement in the classroom, deploying them one-on-one with students may provide more insights about how to make such a teaching approach more engaging.

4.4.3 Effect on Student Engagement in Tablet-Based Education (Region-2)

Next to listening to stories or songs, rote repetition activities are the easiest in our learning application. Although students provided responses rather than passively listening, they only had to repeat what they heard. This resulted in very high scores in rote repetition activities – students scored an average of 97% on all activities regardless of word length or familiarity. However, these high scores did not necessarily lead to the highest student engagement, disputing findings that easy activities are always highly engaging [21]. As reported in the classroom observations in Region-1, we found mixed results on the student engagement with rote repetition activities. This was also evident in the students' usage logs from Region-2. Students had a mean completion rate of 79% on all rote repetition activities, and the addition of real-world photographs did not improve those completion rates.

4.5 Emphasis on Syllables as the unit of reading

4.5.1 Classroom Observation (Region-1)

While pretesting children in Region-1, we found that several students could read simple words and copy simple sentences but could not write the individual letters when asked. We probed further with the help of a native speaker who sang the Swahili alphabet song and asked the children to write down the letters. Rather than write the letters individually, the children sometimes wrote the letters syllabically (see Figure 7), indicating that writing individual letter names was not commonly taught in schools. We also noticed the ABC posters in both schools and determined that students learned their letters in the context of syllable sounds rather than letter names. Given the nature of the Swahili language, this finding was not completely surprising. In contrast to English, Swahili (the national language in Tanzania) is a highly-phonetic language, therefore teaching reading with syllables as the basic unit may be more suitable than letters. Even in English-speaking contexts, prior research shows that children produced more phonologically plausible spellings for nonwords when taught to read with letter sounds, as compared to letter names [10].

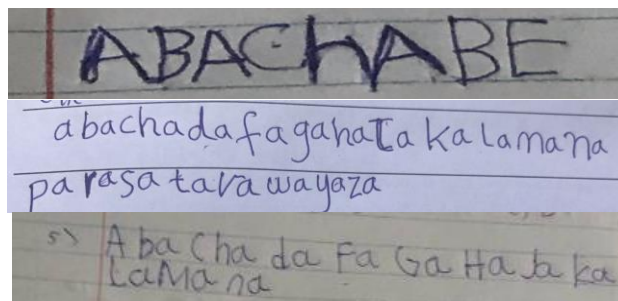


Figure 7: Children's responses to writing letters from A-Z

4.5.2. Design Implications

Following these observations and upon confirmation with the teachers, we included modules in our reading and writing applications to introduce students to vowels and their sounds, then briefly to all letters, and finally teaching them to read and write syllables prior to introducing them to words.

4.5.3 Effect on Student Engagement in Tablet-Based Education (Region-2)

We reported in the baseline testing with children in Region-1, that we found children who could read simple words, but could not name letters. Therefore, we investigated whether this general familiarity with syllables translated to increased student engagement on syllable-based activities in our learning application, even for children without formal schooling experience as in Region-2. Using a paired t-test, we compared the differences between student engagement with all syllable reading activities compared to letter reading activities. We found that students were significantly more engaged with syllables than they were with letters (*mean completion* = 74% vs 51%; *SD* = 26% vs 14%; $t = 8.5701, df = 122, p < 0.01$).

We considered the possibility that the number of questions in a letter activity (range of 10 – 24 questions) was much more than the number of questions in a syllable activity (10 questions each), and that this difference was responsible for the huge variation in completion rates. Therefore, we decided to only compare activities with less than 10 questions to investigate if the different completion rates are still present even with the same number of questions. We found that even after adjusting for question count, students were still more engaged with syllable reading activities via high completion rates than letter reading activities (*mean completion* = 85% vs 76%; *SD* = 18% vs 18%; $t = 3.544, df = 72, p < 0.01$). Although the limited question set increased the completion rates for letter reading activities, our results confirm that in a Swahili-speaking context, students are more engaged when they are taught to read using syllables than when they are with letters.

4.6 Activity difficulty does not reliably predict student engagement

Prior studies have provided mixed results on the correlation between activity difficulty and student engagement. Some studies suggest that students are most engaged in the easiest

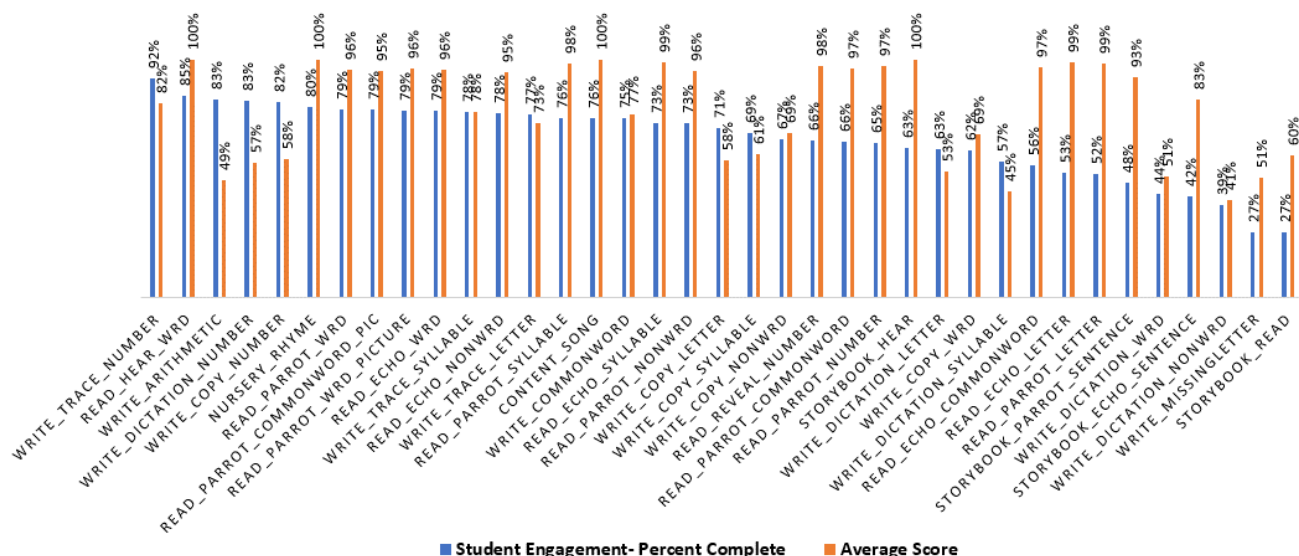


Figure 8: Comparison between student engagement in activity types vs activity difficulty – sorted from highest to lowest on student engagement

activities [21], while others suggest that moderately difficult activities show the most engagement [1,13]. Unlike these studies, we did not find that activity difficulty always correlated with activity engagement. For example, the most engaging activity for the students was tracing numbers (92% average completion rate), even though it was not the easiest activity of the reading and writing applications. On the other hand, several activities such as rote repetition of letters had very high average scores (99%) but very low engagement rates (52% average completion rates - see Figure 8). These findings suggest activity difficulty is not enough to explain student engagement, and there are other factors responsible for explaining engagement, at least in this cultural context.

4.6.1 Prior familiarity with domain area may lead to increased student engagement

Overall, all the number writing activities seemed to be highly engaging for students, including when these activities were challenging. For example, children completed 83% of the activities available for writing arithmetic answers and dictated numbers, despite only responding correctly to just 57% and 49% of such questions, respectively. To explore why number writing activities are so engaging for the students, we explored two theories. First, we hypothesized that the scaffolding provided in the writing application (progress indication, demo activities, and shape manipulation) was motivating, and encouraged the students to persist through the application. However, students did not show the same level of engagement with other writing activities (e.g., letters and syllables) which provided similar feedback. Average scores do not explain these differences either – Figure 8 shows that students’ average score on the writing arithmetic activity was similar to the writing missing letter activity (49% vs 51%), however, the difference in

percent completion of both activities were quite stark (81% vs 28%). Furthermore, a paired t-test confirmed no significant differences in overall student engagement between the reading applications ($M = 66\%$, $SD = 12\%$) and the writing applications ($M = 68\%$; $SD = 21\%$); $t(230) = 1.46$, $p = 0.14$. Therefore, increased student engagement was not a result of the application mechanics.

We also explored the possibility that the increased number writing completion rates were due to students’ higher number prior knowledge. We ran a linear regression model to determine the effects of students’ baseline math scores on their percent completion of the number writing activities. We found that on the Number TRACE activities, students persist for 0.6 % longer for each additional point they scored on the number recognition baseline test (*intercept estimate* 91%; $p = 0.02$). For the Number COPY activity, students persisted for 1% longer for each additional point they scored on the pretest (*intercept estimate* 79%; $p = 0.03$). We found no significant effects of the baseline scores to their engagement in Number DICT and Number ARITH. One explanation is that since student prior knowledge increases engagement in Number TRACE and Number COPY, the similarity of the characters (evidenced by the pretest scores), may encourage them to also persist through Number DICT and ARITH despite the increased difficulty.

Additionally, we confirmed a domain-familiarity effect on overall student engagement; students had a significantly higher engagement with numeracy activities ($M = 75\%$, $SD = 17\%$) vs literacy activities ($M = 66\%$, $SD = 9\%$), $t(226) = 7.51$, $p < 0.01$. This confirms our hypothesis that students’ increased prior familiarity with numbers is

correlated with higher engagement in the numeracy activities, and that this engagement is most evident in the number writing applications.

5 Discussion and Conclusions

Our results show that pedagogical techniques that are engaging in the classroom are potentially also engaging on tablet-based education activities even for children without formal schooling or adequate digital literacy. Given that the logistics of observing and co-creating applications with our target population may be particularly challenging due to access, we show that observing a much smaller set of classrooms can produce insights that increase the engagement of applications on a much larger scale. Previous studies have shown that observing children play, and their daily interaction behaviors, can lead to the design of highly engaging mobile applications [18] – we add to this body of knowledge by showing that classroom observations can provide beneficial insights to create and modify applications that are highly engaging in rural, low-resource contexts. We also show that by using our reading and writing applications for only 6 sessions a month, students demonstrated significant growth in their reading and writing abilities. As researchers investigate other ways to increase student engagement with tablet applications for students in this cultural context, there is the potential for even more learning gains, without the need for students to set foot in a traditional classroom.

Unlike [1,13,20], we found that activity difficulty did not necessarily correlate with student engagement. Prior domain familiarity helped our students persist longer in number writing activities, even through the difficult ones. Our findings reveal that understanding students' prior knowledge does not only give us information on how to appropriately place students within a curriculum, but may also be used as a means to improve engagement throughout an application by purposefully arranging activities and content menus to present students with more engaging options.

Of the pedagogical techniques observed, music was highly engaging in the classroom, as well as on the tablet-based application. This was expected given the high value placed on music as a teaching tool in this cultural context, as well as the plethora of academic, psychological, and cognitive benefits that music brings to learning [12,25]. Furthermore, introducing familiar classroom music through an unfamiliar medium (the tablet), may provide a sense of comfort, and invite collaboration and support from their school-enrolled peers who routinely enjoy these songs in school.

Providing the students with syllable-based activities to support their reading proved to be highly engaging. Although we were not surprised by observing this teaching method given the phonetic nature of Swahili, we were surprised that these effects were evident even with children who have no formal schooling. It shows that there may not only be

academic benefits, but also motivational benefits to presenting students with word roots that make sense in the context of their native language. Our study shows that these students were motivated to learn how to read – it just made more sense to do it using syllables rather than individual letters in this language and cultural context. Research studies show that teaching students using phonetics allows them to produce more plausible spellings to nonwords in English [10] – these effects will be multiplied in a language like Swahili, in which there is a very close mapping of syllable sounds to their written form. This finding does not necessarily suggest that teaching children to read using letters should be eliminated entirely in this cultural context. In Tanzania specifically, children are expected to complete their elementary education in Swahili but transition to English in secondary school to allow for better economic opportunities. The private school we observed was created to help children become better prepared for attending high school in English, as the attrition rates in Region-1 from elementary to high school are very high. Teaching using letters (in addition to syllables) may not be the optimal way to teach Swahili reading, but may better prepare students for future opportunities such as learning and working in an English-speaking context.

Relying on word familiarity to increase student engagement only showed positive results with the writing application. The Inverted-U hypothesis, which stipulates that people will often prefer activities that are marginally more difficult than their comfort zone, may explain this effect [21]. Word familiarity may have caused boredom and therefore less engagement with rote-repetition activities but was enough of a motivational boost to help them persist through the writing activities. Although the value of teaching using rote repetition is debated within the learning science community [38], there are some documented benefits including improved recall ability, as well as introducing students to foundational knowledge [41]. Further research is required to explore other ways to make this activity more engaging.

Our study is limited in the following areas: First, we focused on speech recognition and handwriting input techniques – there are likely many more insights to be gained if other input techniques are studied as well. Also, we gleaned for insights from one region and applied it to another region, although geographically, economically, and culturally similar. There may be other subtle differences between the two populations that we were not aware of. Finally, given that we only analyzed tablet log data for Region-2 participants, we do not have insights into the kinds of support they received from other sources. Friends and family may have provided more support with number writing, and numbers in general, which might partially account for some of their increased engagement with the activity as well.

6 Acknowledgements

We wish to thank our collaborators in Tanzania who have continuously welcomed us into their communities and provided valuable insights that have improved our learning system. We thank all the translators and native Swahili speakers who have ensured that we have translated content and context appropriately across the different media forms. We also wish to thank the Global Learning XPRIZE Initiative for providing us with the participants to conduct such a large-scale study. We thank the Jacobs Foundation for their continued support for our research. Finally, we thank everyone who has contributed to this project by providing software solutions, technical support, domain knowledge support and consulting on our learning content. The research reported here was supported in part by the Institute of Education Sciences, U.S. Department of Education, through grant R305B150008 to Carnegie Mellon University. The opinions expressed are those of the authors and do not represent the views of the Institute or the U.S. Department of Education.

7 References

1. Sami Abuhamdeh and Mihaly Csikszentmihalyi. 2012. The importance of challenge for the enjoyment of intrinsically motivated, goal-directed activities. *Personality and Social Psychology Bulletin* 38, 3: 317–330.
2. Erik Andersen, Yun-En Liu, Rich Snider, Roy Szeto, and Zoran Popović. 2011. Placing a value on aesthetics in online casual games. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 1275–1278.
3. Abhijit V. Banerjee, Shawn Cole, Esther Duflo, and Leigh Linden. 2007. Remedying education: Evidence from two randomized experiments in India. *The Quarterly Journal of Economics* 122, 3: 1235–1264.
4. Cynthia Breazeal, Robin Morris, Stephanie Gottwald, Tinsley Galyean, and Maryanne Wolf. 2016. Mobile devices for early literacy intervention and research with global reach. In *Proceedings of the Third (2016) ACM Conference on Learning@ Scale*, 11–20.
5. Juliet Corbin and Anselm Strauss. 2008. Basics of qualitative research: Techniques and procedures for developing grounded theory. *Thousand Oaks*.
6. Andrea DeBruin-Parecki. 2009. Establishing a family literacy program with a focus on interactive reading: The role of research and accountability. *Early Childhood Education Journal* 36, 5: 385–392.
7. Adele Diamond. 2000. Close interrelation of motor development and cognitive development and of the cerebellum and prefrontal cortex. *Child development* 71, 1: 44–56.
8. Laura H. Dinehart. 2015. Handwriting in early childhood education: Current research and future implications. *Journal of Early Childhood Literacy* 15, 1: 97–118.
9. B. Egeland, N. Weinfeld, M. Hiester, C. Lawrence, S. Pierce, K. Chippendale, and J. Powell. 1995. Teaching tasks administration and scoring manual. *University of Minnesota*.
10. Michelle R. Ellefson, Rebecca Treiman, and Brett Kessler. 2009. Learning to label letters by sounds or names: A comparison of England and the United States. *Journal of Experimental Child Psychology* 102, 3: 323–341.
11. Lita Ericson and Moira Fraser Juliebo. 1998. *The Phonological Awareness Handbook for Kindergarten and Primary Teachers*. ERIC.
12. Joyce Eastlund Gromko. 2005. The effect of music instruction on phonemic awareness in beginning readers. *Journal of research in music education* 53, 3: 199–209.
13. Susan Harter. 1978. Pleasure derived from challenge and the effects of receiving grades on children’s difficulty level choices. *Child Development*: 788–799.
14. David Huggins-Daines, Mohit Kumar, Arthur Chan, Alan W. Black, Mosur Ravishankar, and Alexander I. Rudnicky. 2006. Pocketsphinx: A free, real-time continuous speech recognition system for hand-held devices. In *2006 IEEE International Conference on Acoustics Speech and Signal Processing Proceedings*, I–I.
15. Karin H. James and Laura Engelhardt. 2012. The effects of handwriting experience on functional brain development in pre-literate children. *Trends in neuroscience and education* 1, 1: 32–42.
16. Jignesh Khakhar and Sriganesh Madhvanath. 2010. Jollymate: Assistive technology for young children with dyslexia. In *2010 12th International Conference on Frontiers in Handwriting Recognition*, 576–580.
17. Anuj Kumar, Pooja Reddy, Anuj Tewari, Rajat Agrawal, and Matthew Kam. 2012. Improving literacy in developing countries using speech recognition-supported games on mobile devices. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 1149–1158.
18. Anuj Kumar, Anuj Tewari, Geeta Shroff, Deepti Chittamuru, Matthew Kam, and John Canny. 2010. An exploratory study of unsupervised mobile learning in rural India. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 743–752.
19. Marlaine E. Lockheed and Adriaan M. Verspoor. 1991. *Improving primary education in developing countries*. Oxford University Press for World Bank.
20. Derek Lomas, Dixie Ching, Eliane Stampfer, Melanie Sandoval, and Ken Koedinger. 2011. “ Battleship Numberline”: A Digital Game for Improving Estimation Accuracy on Fraction Number Lines. *Society for Research on Educational Effectiveness*.
21. Derek Lomas, Kishan Patel, Jodi L. Forlizzi, and Kenneth R. Koedinger. 2013. Optimizing challenge in an educational game using large-scale design experiments. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 89–98.
22. Zupai Luo, Paul E. Jose, Carol S. Huntsinger, and Therese D. Pigott. 2007. Fine motor skills and

- mathematics achievement in East Asian American and European American kindergartners and first graders. *British Journal of Developmental Psychology* 25, 4: 595–614.
23. G. R. Lyon. 1997. *How do children learn to read?: Statement before the Committee on Education and the Workforce, US House of Representatives*. Washington, DC.
 24. Michael A. Madaio, Fabrice Tanoh, Axel Blahoua Seri, Kaja Jasinska, and Amy Ogan. 2019. Everyone Brings Their Grain of Salt: Designing for Low-Literate Parental Engagement with a Mobile Literacy Technology in Côte d'Ivoire. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, 465.
 25. Allan Miller and Dorita Coen. 1994. The case for music in the schools. *The Phi Delta Kappan* 75, 6: 459–461.
 26. G. Ayorkor Mills-Tettey, Jack Mostow, M. Bernardine Dias, Tracy Morrison Sweet, Sarah M. Belousov, M. Frederick Dias, and Haijun Gong. 2009. Improving child literacy in Africa: experiments with an automated reading tutor. In *Information and Communication Technologies and Development (ICTD), 2009 International Conference on*, 129–138.
 27. Sugata Mitra, Ritu Dangwal, Shiffon Chatterjee, Swati Jha, Ravinder S. Bisht, and Preeti Kapur. 2005. Acquisition of computing literacy on shared public computers: Children and the "hole in the wall". *Australasian Journal of Educational Technology* 21, 3.
 28. Amelia K. Moody, Laura M. Justice, and Sonia Q. Cabell. 2010. Electronic versus traditional storybooks: Relative influence on preschool children's engagement and communication. *Journal of Early Childhood Literacy* 10, 3: 294–313.
 29. David W. Moore, Thomas W. Bean, Deanna Birdyshaw, and James A. Rycik. 1999. Adolescent literacy: A position statement. *Journal of Adolescent & Adult Literacy* 43, 1: 97–112.
 30. Aidan Mulkeen. 2005. *Teachers for rural schools: A challenge for Africa*. Africa region World Bank. Washington, DC: World Bank.
 31. Tony Read. 2015. *Where have all the textbooks gone?: Toward sustainable provision of teaching and learning materials in Sub-Saharan Africa*. World Bank Publications.
 32. Kathleen Roskos, Karen Burstein, and Byeong-Keun You. 2012. A Typology for Observing Children's Engagement with eBooks at Preschool. *Journal of Interactive Online Learning* 11, 2.
 33. Yusuf Sayed. 2012. Making education a priority in the Post-2015 development agenda: report of the Global Thematic Consultation on Education in the Post-2015 Development Agenda - UNESCO Digital Library. Retrieved March 10, 2019 from <https://unesdoc.unesco.org/ark:/48223/pf0000223024>
 34. John A. Smith. 2000. Singing and songwriting support early literacy instruction. *The reading teacher* 53, 8: 646–649.
 35. Catherine E. Snow and Connie Juel. 2005. Teaching children to read: What do we know about how to do it?
 36. Seung-Hee Son and Samuel J. Meisels. 2006. The relationship of young children's motor skills to later reading and math achievement. *Merrill-Palmer Quarterly* (1982-): 755–778.
 37. Mariama Sow. 2017. Figures of the week: Africa, education, and the 2018 World Development Report. *Brookings*. Retrieved November 4, 2018 from <https://www.brookings.edu/blog/africa-in-focus/2017/10/06/figures-of-the-week-africa-education-world-development-report-2018/>
 38. Paweł Szudarski and Kathy Conklin. 2014. Short-and Long-Term Effects of Rote Rehearsal on ESL Learners' Processing of L2 Collocations. *TESOL quarterly* 48, 4: 833–842.
 39. Judith Uchidiuno, Evelyn Yarzebinski, Michael Madaio, Nupur Maheshwari, Ken Koedinger, and Amy Ogan. 2018. Designing Appropriate Learning Technologies for School vs Home Settings in Tanzanian Rural Villages. In *Proceedings of the 1st ACM SIGCAS Conference on Computing and Sustainable Societies*, 9.
 40. Tanvi Zunjarrao and Uday Joshi. 2017. Recognition of Online Handwritten Characters Using LIPI Toolkit. *International Journal of Advance Research, Ideas and Innovations in Technology* 3: 689–692.
 41. 2017. Rote Learning vs. Meaningful Learning. *Oxford Learning*. Retrieved March 12, 2019 from <https://www.oxfordlearning.com/difference-rote-learning-meaningful-learning/>